

REMARKS

Claims 1-20 are pending. Independent claims 1, 11, and 20 were rejected under 35 U.S.C. 102(b) as being anticipated by Sharma (U.S. Patent No. 6,108,737). The Examiner also rejected the independent claims as being unpatentable under 35 U.S.C. 103(a) as being unpatentable over Hum (U.S. Patent Application Publication No. 2004/0123047) in view of Bauman (U.S. Patent No. 6,189,078).

Sharma 35 U.S.C. 102(b) Rejection

The Examiner provided a first basis for rejecting the independent claims by citing Sharma. The Examiner argues that Sharma teaches a cluster of processors or a plurality of processors interconnected using point-to-point architecture. The Applicants respectfully disagree. Sharma describes processors that are connected by a "switch 200." However, the "switch 200" presented in Sharma does not connect the processors in a point-to-point architecture, but instead merely routes "memory reference operations ... to the Arb bus 170" (Sharma 6, 60-65). The processors can not communicate with each other using point-to-point links but instead must have operations processed by "an arbiter 240" that "arbitrates among the input queues to grant access to the Arb bus 170 where the requests are ordered into a memory reference request stream" (Sharma 7, 53-55). Sharma uses a conventional I/O bus that "may operate according to the conventional Peripheral Computer Interconnect (PCI) protocol" (Sharma 5, 62-63). A conventional bus such as PCI is a shared bus architecture that is distinct from the point-to-point architecture recited in all the independent claims. As noted in the present application, "In a point-to-point architecture, a cluster of processors includes multiple processors directly connected to each other through point-to-point links. By using point-to-point links instead of a conventional shared bus or external network, multiple processors are used efficiently in a system sharing the same memory space." (Application page 5, lines 22-25)

The Examiner further argues that Sharma describes a remote data cache recited in the independent claims. The Examiner argues that a remote data cache is "one of the other private caches 122-128 of the processor." The Applicants respectfully disagree. The present application explicitly states that one of the other private caches is not a remote data cache. As noted in the present application, "Unlike the caches associated with each processor, in this example a single Application No.: 10/635,703

remote data cache is provided for multiple processors in a cluster. If a memory line requested by a particular processor is not found in the cache associated with the processor, a cache coherence controller does not necessarily need to forward the request to a home cluster. Instead the cache coherence controller can check the remote data cache to determine if the memory line is resident in the remote data cache. If the memory line is resident in the remote data cache, no request needs to be forwarded to the home cluster and the cache coherence controller can instead use the data in the remote data cache to respond to the transaction” (Application page 21, lines 26 – page 22, line 2). A remote data cache is a mechanism for “reducing intercluster transactions” (Application page 21, line 20). “Any cache holding data from remote clusters accessed by local processors that is accessed after local caches but before a home cluster memory controller is referred to herein as a remote data cache” (Application page 22, line 14-18). None of the references cited including Sharma teach or suggest such an entity that is accessed “after local caches but before a home cluster memory controller.”

Sharma does not teach or suggest any remote data cache and in fact does not even make reference to reducing intercluster transactions. For at least these reasons, it is respectfully submitted that the independent claims and associated dependent claims are allowable over Sharma.

Hum/Bauman 35 U.S.C. 103(a)/102(e) Rejection

In the previous office action response, the Applicants argued that Hum does not describe “providing response information with a completion indicator to the processor” as recited in independent claims 1, 11, and 20 and that the claims were allowable over Hum for at least these reasons. The Applicants also amended the claims to facilitate prosecution. However, the Applicants would like to emphasize that Hum does not teach a variety of elements recited in the independent claims, even when Hum is combined with Bauman. The Examiner responded by changing the anticipation rejection to an obviousness rejection and by arguing that completion responses are well-known in the cache-coherency art as evidenced by description in Bauman. The Applicants respectfully disagree.

The Examiner argues that Hum describes a processing cluster connected using a
Application No.: 10/635,703

point-to-point architecture. The Applicants respectfully disagree. Hum does have multiple processors, but the processors are connected using a shared bus 710 (Hum Fig. 7). Hum makes no reference to any point-to-point architecture.

The Examiner also argues that Hum teaches a remote data cache in its description or an import cache 250. The Applicants respectfully disagree. The import cache 250 "can also avoid broadcasting requests from elsewhere in the system to the agent's local cluster. The agent can use the import cache to determine that no nodes within the cluster have a copy of the cache line" (Hum paragraph 65) A remote data cache is not used to determine that no nodes within the cluster have a copy of the cache line. By contrast, a remote data cache holds data from remote clusters, not from within the cluster. "Any cache holding data from remote clusters accessed by local processors that is accessed after local caches but before a home cluster memory controller is referred to herein as a remote data cache" (Application page 22, line 14-18)

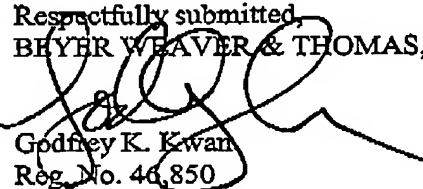
Although Hum does not teach or suggest a completion response, the Examiner argues that Bauman determining that the cache access request can be handled locally by using the remote data cache without having to probe remote nodes. As noted above, the remote data cache is described in the specification as an entity that holds data from nodes in remote clusters. Bauman does not teach or suggest any such entity and consequently can not provide a "completion indicator to the processor when it is determined that the cache access request can be handled locally" as is recited in the claims. The material the Examiner cites merely describes "Memory Storage Unit" or MSU responses. The MSU is not a remote data cache. The acknowledgments are merely responses from an MSU and not completion indicators provided when it is determined that "the cache access request can be handled locally by using the remote data cache without having to probe remote nodes" as is recited in the claims.

Even if it were appropriate to combine Bauman and Hum, Hum does not teach or suggest a point-to-point architecture and Bauman actually teaches away from processor clusters or a plurality of processors interconnected using a point-to-point architecture as is explicitly recited in the independent claims. Bauman actually describes a crossbar switch as an inefficient system and actually teaches away from using a point-to-point architecture such

Application No.: 10/635,703

as a crossbar switch. A crossbar system "may not be readily expanded. A single crossbar switching network has a predetermined number of switched cross points placed at intersections between the processors and memory modules." (Bauman col. 2, lines 54-57) Bauman continues to note that "a single crossbar switching network has a predetermined number of switched cross points placed at intersections between the processors and memory modules. These switched cross points may accommodate a predetermined maximum number of processors and memory modules. Once each of the switched cross points is utilized, the system may not be further expanded." (Bauman col. 2, lines 58-63) Moreover, such a distributed system poses an increased challenge for maintaining memory coherency. Although an invalidation approach similar to the one described above may be utilized, the routing of these requests over each of the point-to-point interfaces to each of the local caches associated with the processors increases system overhead." (Bauman col. 2, lines 58-67)

In light of the above remarks relating to independent claims and certain dependent claims, the remaining dependent claims are believed allowable for at least the reasons noted above. Applicants believe that all pending claims are allowable and respectfully request a Notice of Allowance for this application from the Examiner. Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

Respectfully submitted,
BHYER WEAVER & THOMAS, LLP

Godfrey K. Kwan
Reg. No. 46,850

P.O. Box 70250
Oakland, CA 94612-0250
(510) 663-1100

Application No.: 10/635,703

-7-